ORIGINAL

KIRKLAND & ELLIS

PARTNERSHIPS INCLUDING PROFESSIONAL CORPORATIONS

EX PARTE OR LATE FILED

655 Fifteenth Street, N.W.
Washington, D.C. 20005

202 879-5000

Facsimile:
202 879-5200

John P. Frantz
To Call Writer Directly:
(202) 879-5038
john_frantz@kirkland.com

November 10, 1999

BY HAND DELIVERY

Ms. Magalie Roman Salas
Secretary, Federal Communications Commission
455 12th Street, S.W.
Room TW B204
Washington, D.C. 20554

Re: *AT&T/MediaOne Merger*, CS Docket No. 99-251

Dear Ms. Salas:

Enclosed please find an original and four copies of the Written Ex Parte of Professor Mark A. Lemley and Professor Lawrence Lessig. Pursuant to the Commission's Public Notice dated August 23, 1999, please accept this document as part of the public record in the above referenced proceeding.

Sincerely,

John P. Frantz

Enclosure

cc: To-Quyen Truong
Frances Eisenstein
Walter Strack
Sunil Daluvoy
International Transcription Service, Inc.

No. of Copies rec'd 0+4
List ABCDE

CERTIFICATE OF SERVICE

I, John P. Frantz, certify that two copies of the enclosed Written Ex Parte of Professor Mark A. Lemley and Professor Lawrence Lessig were served by first-class mail, postage prepaid, on the following:

Mark C. Rosenblum
Steven C. Garavito
Lawrence J. Lafaro
AT&T CORP.
Room 3252G1
295 North Maple Avenue
Basking Ridge, NJ 07920

Susan M. Eid
Sean C. Lindsay
MEDIAONE GROUP, INC.
1919 Pennsylvania Avenue, N.W.
Suite 610
Washington, D.C. 20006

Howard J. Symons
Michael M. Mundt
MINTZ, LEVIN, COHN, FERRIS, GLOVSKY & POPEO, P.C.
700 Pennsylvania Avenue, N.W.
Suite 900
Washington, D.C. 20004

Wesley R. Heppler
Robert L. James
COLE, RAYWID & BRAVERMAN, L.L.P.
1919 Pennsylvania Avenue, N.W.
Suite 200
Washington, D.C. 20006

David W. Carpenter
Mark D. Schneider
David L. Lawson
Lorrie M. Marcil
C. Frederick Beckner
SIDLEY & AUSTIN
1722 Eye Street, N.W.
Washington, D.C. 20006

Phillip L. Verveer
Michael H. Hammer
Michael G. Jones
Francis M. Buono
WILLKIE, FARR & GALLAGHER
1155 Twenty-First Street, N.W.
Suite 600
Washington, D.C. 20036

John P. Frantz

November 10, 1999

ORIGINAL

FX PARTE OR LATE FILED

Before the

FEDERAL COMMUNICATIONS COMMISSION

Washington, D.C. 20554

In the Matter of:	
Application for Consent to the Transfer of Control of Licenses MediaOne Group, Inc. to AT&T Corp.	CS Docket No. 99-251

WRITTEN EX PARTE OF PROFESSOR MARK A. LEMLEY AND PROFESSOR LAWRENCE LESSIG

INTRODUCTION

1. We offer this written ex parte to address the question of "open access" and its relationship to the architecture of the Internet. It is our view that the extraordinary growth and innovation of the Internet depends crucially upon this architecture. Changes in this architecture should be viewed with skepticism, as they may in turn threaten this innovation and growth.

2. The proposed design of the merged AT&T and MediaOne entity, once allied with the existing merged AT&T and TCI, threatens to compromise an important architectural principle of the Internet. As we describe in more detail below, it threatens to weaken the Internet's "End-to-End" design. In our view, this change could have profound implications for the future of growth and innovation on the net.

3. The FCC's analysis to date does not consider these principles of the Internet's design. It therefore does not adequately evaluate the potential threat that this merger presents. Neither does the FCC's approach properly account for its role in creating the conditions that

made the Internet possible. Under the banner of "no regulation," the FCC threatens to permit this network to calcify as earlier telecommunications networks did. Further, and ironically, the FCC's supposed "hands off" approach will ultimately lead to more rather than less regulation.

4. We do not yet know enough about the relationship between these architectural principles and the innovation of the Internet. But we should know enough to be skeptical of changes in its design. The strong presumption should be in favor of preserving the architectural features that have produced this extraordinary innovation. The FCC's presumption should be against approving mergers that threaten these design principles, without a clear showing that the threat would not undermine the Internet's innovation. No such showing has been made in this case.

5. In Part I of this declaration, we explain our background and interest in this matter. In Part II, we describe the design principles of the Internet, and how they differ from the principles animating traditional telephone networks. In Part III, we explain why permitting AT&T to bundle ISP service with network access threatens the structure of the Internet. Finally, in Part IV we respond to arguments that have been made to permit the merged AT&T/MediaOne to extend its monopoly to control ISP service.

I. Background

6. Lemley is the Marrs McLean Professor of Law at the University of Texas in Austin, Texas, where he teaches intellectual property, computer law, patent law, antitrust, electronic commerce and regulation of the Internet. Beginning January 1, 2000 he will assume an appointment as Professor of Law at the Boalt Hall School of Law, University of California at Berkeley. He is of counsel to the law firm of Fish & Richardson, where he litigates and counsels clients in the areas of antitrust, intellectual property and computer law. He is the author of four

books and twenty-eight articles on these and related subjects, has taught intellectual property law to federal judges at the Federal Judicial Center, and has testified before Congress and the Federal Trade Commission on patent and antitrust matters. His articles have appeared in the Yale Law Journal, the Stanford Law Review, the California Law Review, the Texas Law Review, the Duke Law Journal, and the Southern California Law Review, as well as numerous specialty journals. He has chaired or co-chaired a dozen major conferences on intellectual property and computer law, including Computers Freedom and Privacy '98, and he was the 1997 Chair of the Association of American Law Schools Section on Law and Computers. He received his J.D. from Boalt Hall School of Law at the University of California at Berkeley, and his A.B. from Stanford University. After graduating from law school, he clerked for Judge Dorothy Nelson on the United States Court of Appeals for the Ninth Circuit, and practiced law in Silicon Valley with Brown & Bain and with Fish & Richardson before coming to Texas.

7. Lessig is the Jack N. and Lillian R. Berkman Professor for Entrepreneurial Legal Studies at the Harvard Law School, where he teaches courses related to the law of cyberspace. He has just completed a book, *Code and Other Laws of Cyberspace*, which analyzes the relationship between the architecture of the Internet and the freedoms the Internet enables. Lessig has written many articles, both for scholarly and popular journals, about the Internet and its regulation. Among scholarly journals, he has published articles in the Yale Law Journal, Stanford Law Journal, Harvard Law Review (forthcoming), Michigan Law Journal (forthcoming), Emory Law Journal, and the Proceedings of the IEEE (forthcoming); among popular journals, he has published articles in The Industry Standard, The New Republic, and Wired Magazine. Lessig teaches constitutional law, contracts, comparative constitutional law, and the law of cyberspace. In 1995, he taught the basic antitrust course at the Yale Law School,

as well as a course on the law of cyberspace. Except for this year, he has continued to teach courses related to the Internet and its regulation, including a seminar examining *United States v. Microsoft*. He graduated from Yale Law School in 1989, and clerked for Judge Richard A. Posner of the Seventh Circuit Court of Appeals, and then Justice Antonin Scalia of the Supreme Court. In 1991, he was appointed to the faculty of the University of Chicago Law School, where he received tenure in 1994. He was a visiting professor at Yale Law School in 1995. During 1996-97, he was a fellow at the Program for Ethics and the Professions at Harvard University. In 1997, he joined the Harvard Law School. He is currently a fellow at the Institute for Advanced Study in Berlin.

8. We are not experts in computer architecture, or in software technology. We have studied these subjects to understand their relationship to the competitive environment that the Internet has created, and to the values that we have come to recognize in the Internet. It is our view that any analysis of policy related to the Internet must explicitly consider these architectural aspects of the Internet's design. This "architecture," as the Electronic Frontier Foundation puts it, "is policy." Changing the architecture is to change that policy.

9. We have not been retained by any party in this matter. At the request of Bell Atlantic and GTE, Lessig attended a meeting with the Justice Department, to discuss the AT&T/MediaOne merger, and helped arrange a meeting with Bell Atlantic and GTE, and some in the Internet community to discuss broadband cable access. Bell Atlantic and GTE paid for Lessig's expenses relating to those two meetings. Lemley has had no involvement with any party with an interest in this litigation.

10. This declaration does not address the general question of the merger of AT&T and MediaOne beyond its effects on the market for broadband Internet service. We also assume

for purposes of this declaration that "residential broadband access" is properly considered a separate antitrust market. Our statements here relate solely to the question of the architecture of the resulting network under that assumption.

11. It is our view that it is important for the FCC to consider these matters now in the context of this merger. It is not our view that every entity that connects its network to the Internet must obey the principles that we describe below. Innovation in the Internet generally is not threatened by what any one (small) company on the Internet might do. But when a policy of closed access covers a portion of the net as significant as AT&T's would after this merger, the FCC must consider the overall effect that AT&T's proposed change in architecture would have on innovation. More importantly, even if AT&T's policy affected the same footprint of the net before the merger as after the merger, the increased ability post merger for the actors to behave strategically increases the significance of the threat. So far, the FCC has considered only the effect its "regulations" would have on broadband investment. In our view, that is only one part of the equation. The more significant part is the effect that the FCC's *failure to regulate* will have on innovation generally.[1]

II. Design Principles of the Internet

A. The Architecture of the Internet

12. The Internet is the fastest growing network in history. In the 30 years of its life, its population has grown a million times over. It is currently the single largest contributor to the growth of the United States economy, and has become the single most important influence linking individuals, and commerce, internationally.

---

[1] For an analysis similar to our own, see François Bar, Stephen Cohen, Peter Cowhey, Brad DeLong, Michael Kleeman, John Zysman, *Defending the Internet Revolution in the Broadband Era: When Doing Nothing is Doing Harm*, E-conomy Working Paper 12, Berkeley Roundtable on the International Economy (BRIE), August 1999.

13. It is not, however, the first communications network. There have been other networks before the Internet that did not experience the same extraordinary growth. These networks followed different design principles, including different principles governing how protocols would evolve and become adopted. It is our view that these differences in growth can be traced, at least in part, to these differences in design.

14. It is a view of many in the Internet community, and ours as well, that the extraordinary growth of the Internet rests fundamentally upon its design principles. Some of these principles relate to the openness of the Internet's standards and the openness of the software that implemented these standards. Some are engineering principles, designed to make the net function more flexibly and efficiently. But from the very beginning, these principles have been understood to have a social as well as technological significance. They have, that is, been meant to implement values as well as enable communication. In our view, one aspect of this social significance is the competition in innovation the Internet enables. The tremendous innovation that has occurred on the Internet, in other words, depends crucially on these design principles.

15. Among the Internet's design principles is one that is particularly relevant to these proceedings. This is the "End-to-End" design principle has been latent in system design for many years, but was first articulated explicitly as principle in 1981 by Professors Jerome H. Saltzer, David P. Reed, and David D. Clark.[2]

16. The "End-to-End" principle organizes the placement of functions within a network. It counsels that that "intelligence" in a network be located at the top of a layered system

---

[2] See *End to End Arguments in System Design*, http://web.mit.edu/Saltzer/www/publications/.

— at its "ends," where users put information and applications onto the network — and that the communications protocols themselves (the "pipes" through which information flows) be as simple and general as possible.

17. One consequence of this design is a principle of non-discrimination among applications. Lower-level network layers should provide a broad range of resources that are not particular to or optimized for any single application — even if a more efficient design for at least some applications is thereby sacrificed. As described in a subsequent paper by Reed, Saltzer, and Clark,

> End to end arguments have ... two complimentary goals: (1) Higher-level layers, more specific to an application, are free to (and thus expected to) organize lower level network resources to achieve application-specific design goals efficiently (application autonomy); (2) lower-level layers, which support many independent applications, should provide only resources of broad utility across applications, while providing to applications useable means for effective sharing of resources and resolution of resource conflicts (network transparency).[3]

18. While the End-to-End design principle was first adopted for technical reasons, it has important social and competitive features as well. End-to-end expands the competitive horizon, by enabling a wider variety of applications to connect and use the network. It maximizes the number of entities that can compete for the use and applications of the network. As there is no single strategic actor who can tilt the competitive environment (the network) in favor of itself, or no hierarchical entity that can favor some applications over others, an End-to-End network creates a maximally competitive environment for innovation, which by design assures competitors that they will not confront strategic network behavior.

---

[3] See David P. Reed, Jerome H. Saltzer, and David D. Clark, *Comment on Active Networking and End-to-End Arguments*, IEEE Network 12, 3 (May/June 1998) pages 69-71.

19. The End-to-End design of the Internet has facilitated innovation. As Reed, Saltzer and Clark argue, for example: "had the original Internet design been optimized for telephony-style virtual circuits (as were its contemporaries SNA and TYMNET), it would not have enabled the experimentation that led to protocols that could support the World-Wide Web, or the flexible interconnect that has led to the flowering of a million independent Internet Service providers. *Preserving low-cost options to innovate outside the network, while keeping the core network services and functions simple and cheap, has been shown to have very substantial value.*"[4]

20. The principle of End-to-End is not unique to computer networks. It has important analogs in American constitutional law and in other legal contexts. Vis-à-vis the states, for example, the dormant commerce clause imposes an End-to-End design on the flow of commerce: No state is to exercise a control over the flow of commerce between states; and the kind of control that a state may exercise over commerce flowing into that state is severely limited. The "network" of interstate commerce is to be influenced at its ends — by the consumer and producer — and not by intermediary actors (states) who might interfere with this flow for their own political purposes. Vis-à-vis transportation generally, End-to-End is also how the principle of common carriage works. The carrier is not to exercise power to discriminate in the carriage. So long as the toll is paid, it must accept the carriage that it is offered. In both contexts, the aim is to keep the transportation layer of intercourse simple, so as to enable the multiplication of applications at the end.

---

[4] *Id.* at 70 (emphasis added). Note the initial ARPANET did not implement End-to-End perfectly into its design. It was because of changes in the 1970s suggested by Vint Cerf and David P. Reed that the network we now recognize as the Internet conformed to End-to-End.

B. The Consequences of these Architectural Principles

21. The effect of these Internet design principles — including, but not exclusively, End-to-End — has been profound. By its design, the Internet has enabled an extraordinary creativity precisely because it has pushed creativity to the ends of the network. Rather than relying upon the creativity of a small group of innovators who work for the companies that control the network, the End-to-End design enables anyone with an Internet connection to design and implement a better way to use the Internet. By architecting the network to be neutral among uses, the Internet has created a competitive environment where innovators know that their inventions will be used if useful. By keeping the cost of innovation low, it has encouraged an extraordinary amount of innovation.

22. The contexts in which this innovation has occurred are many. By keeping the network simple, and its interaction general, the Internet has facilitated the design of applications that could not have originally been envisioned. And by keeping the cost of innovation low in the future — especially in the context of broadband media — the Internet should continue to facilitate innovation.

23. End-to-end design does not only promote innovation by creating the opportunity for innovators to offer services to the network. In our view, the effect comes as well from the expectation that innovation will not be countered by strategic actors who might control the flow of commerce. The potential of an actor in the distributional network to act strategically is a cost to innovation. The expectation that an actor can act strategically is an expected cost to innovation. Thus to the extent an actor is structurally capable of acting strategically, the rational innovator will reckon that capacity as a cost of innovation. Compromising End-to-End will, then, tend to undermine innovation.

24. The End-to-End design of the Internet thus minimizes the cost of strategic behavior, while creating an extraordinary market that innovators can rely upon when developing new applications for the Internet.

C. The Difference with the Architectural Principles of the Old Telephone Network

25. The Internet's design principles are different from the design principles that governed the telephone network prior to the series of actions by the FCC and Antitrust Division of the Justice Department that resulted in the break-up of AT&T. Prior to that break-up, the telephone network was not governed by the principles of End-to-End. The old telephone network was not neutral about the uses to which the telephone system could be placed. For much of the history of the telephone network, it was a crime to use the network in ways not specified by the AT&T. It was a crime, for example, to attach devices that performed services not offered by AT&T, or to provide services that competed with the services provided by AT&T. In the 1940s, even the telephone book was owned by AT&T.

26. Innovation under the old design was thus controlled by AT&T. If a person with a competing conception of how a communications network should be designed wanted to implement that competing conception, he or she would have to either work for AT&T, or convince AT&T of the merits of this alternative design. AT&T was, therefore, a bottleneck on creativity in network architecture. While no doubt AT&T did much to advance telecommunications, through Bell Labs and other research, it also decided which innovations would be deployed. No doubt its decision turned in part upon the expected effect a new technology would have on AT&T's own business model.

27. The early history of the Internet was affected by this control. As described by John Naughton in *A Brief History of the Future* (1999), an early design idea for the Internet was proposed to AT&T by RAND researcher, Paul Baran, in the early 1960s. Resistance to his design was strongest from AT&T. As Naughton reports, Baran recalls one particularly telling instance of AT&T's opposition:

> [AT&T's] views were once memorably summarised in an exasperated outburst from AT&T's Jack Osterman after a long discussion with Baran. 'First,' he said, 'it can't possibly work, and if it did, damned if we are going to allow the creation of a competitor to ourselves.'[5]

28. This resistance is perfectly understandable. From AT&T's perspective, maximizing its control over its network was no doubt profit maximizing. And we should expect corporate entities to behave in a profit maximizing manner. But this resistance was profit maximizing only because AT&T was in control of the network uses. Or in other words, only because the network was not "End-to-End." Had the network been End-to-End, it would have had no incentive to disable one use of the network it controlled in favor for another.

29. The same point about the relationship between innovation and the concentration of control can be made more obviously about the Internet in foreign countries. It is no accident that the Internet was born in the United States, since in practically every other nation, the telephone architecture was controlled by state sponsored monopolies. These monopolies, no less than AT&T, had no interest in facilitating the design of a network that would free individuals from that control. For much of the 1990s, it was a crime in parts of Europe to connect a modem

---

[5] **John Naughton, A Brief History of the Future** 107 (1999). Authors Katie Hafner and Matthew Lyon recount a similar resistance in **Where Wizards Stay Up Late** 52-66 (1996).

to a telephone line. Even today, the franchise in Germany for public phones permits the provider to control how access to the Internet occurs.

D. The Government's Role in Creating the Competitive Environment for the Internet

30. It is fashionable today to argue that innovation is assured if government simply stays out of the way. The FCC's hands-off policy to date appears largely to be motivated by this prevailing ideological vogue. The view is that the best way for the government to guarantee growth in Internet broadband is to let the owners of networks architect broadband as they see fit.

31. We believe this view is misguided. It ignores the history that gave the Internet its birth, and threatens to reproduce the calcified network design that characterized our communications network prior to the Internet. The restrictions on innovation that marked the AT&T telephone monopoly were not removed by the government doing nothing. They were removed by active intervention designed to assure the possibility for innovation. It was the FCC and Department of Justice that cut the knot that tied innovation on the telecommunications network to the innovation favored by AT&T. It was their action that eventually freed the network from the control of a single strategic actor, and opened it up for the innovation of many.

32. Beginning with the Carterfone decision in 1968,[6] the FCC increasingly pursued a policy that forced AT&T to open its network to competing uses and providers. In a series of decisions, the FCC required that AT&T permit alternative uses of its network. In 1984, actions by the Antitrust Division forced AT&T to unbundle its long-distance service from its local telephone service. This unbundling was effected through a decree that lead to the breakup of the largest monopoly in American history.

33. These actions together transformed the telephone network from a network whose use was controlled by one company — AT&T — into a general purpose network, whose ultimate use was determined by end users. In effect, they imposed a principle analogous to End-to-End design on the telephone network. Indeed, though it masquerades under a different name ("open access"), this design principle is part and parcel of recent efforts by Congress and the FCC to deregulate telephony. The fundamental economic goal of the FCC in deregulating telephony is to isolate the natural monopoly component of a network — the actual wires — from other components in which competition can occur. By requiring the natural monopoly component at the basic network level to be open to competitors at higher levels, intelligent regulation can minimize the economic disruption caused by that natural monopoly, and permit as much competition as industry structure will allow.

34. It is our view that but for these changes brought about by the government, the Internet as we know it would not have been possible. Without these changes, the trend in telecommunications was towards more centralized control over the communication network. Network theorist Robert Fano of MIT, for example, wrote in 1972 that unless there was a change in the trend in the computer-communications network, existing institutions would further isolate computer and communications technologies from broad based control.[7] But by seeding the development of a network within a different communication paradigm, and then opening the existing communication network so that it might deploy this different communication paradigm, the government created the conditions for the innovation that the Internet has realized.

---

[6] In the Matter of Use Of The Carterfone Device In Message Toll Telephone Service; Docket No. 16942; 13 F.C.C.2d 420; June 26, 1968.

[7] See, Robert M. Fano, *On The Social Role of Computer Communications*, 60 **Proceedings of the IEEE,** September 1249 (1972).

35. This is not to say that the government created the innovation that the Internet has enjoyed. Nor is it to endorse government, rather than private, development of Internet-related technologies. Obviously, the extraordinary innovation of the Internet arises from the creativity of private actors from around the world. Some of these actors work within corporations. Some of the most important have been associated with the Free Software, and Open Source Software Movements. And some have been entrepreneurs operating outside of any specific structure. But the creativity that these innovators have produced would not have been enabled but for the opening of the communications network. Our only point is that the government had a significant role in that opening.

36. We do not claim that no communication network would have been possible without the government's intervention. Obviously, we have had telecommunication networks for over a hundred years; and as computers matured, we no doubt would have had more sophisticated communication-computer networks. But the design of those networks would not have been the design of Internet. The design would have been more like the French "equivalent" to the Internet — miniTel. But miniTel is not the Internet. The miniTel is a corporate, centralized, controlled version of the Internet. And it is notably less successful.

E. The Relevance of Legacy Monopolies

37. As we have said, no one fully understands the dynamics that have made the innovation of the Internet possible. But we do have some clues. One important element of that innovation is a structure that disables the power of legacy monopolies to influence the future of a network design.

38. By freeing the telecommunications network from the control of one single actor, the government enabled innovation free from the influences of what one might call "legacy"

business models. Companies develop core competencies, and most of them tend to stick to what they know how to do. Companies faced with a potential for radical change in the nature of their market may recoil, either because they don't know how to change to face changing conditions, or because they fear that they will lose the dominance they had in the old market as it becomes a new playing field. Their business planning is, in short, governed by the legacy of their past success. These legacy business plans often affect a company's plans about how to respond to innovation. In a competitive environment, they will often disadvantage a company that fails to respond rapidly enough to changed circumstances.

39. In some markets, companies have no choice but to respond to changed circumstances. They either change or die. It is a mark of Microsoft's success, for example, that its chairman, Bill Gates, succeeded in radically altering the course of the company's development, in the face of changed competitive circumstances, despite the fact that such changes resulted in the termination of projects at other times deemed central to Microsoft's future (MSN, for example.). In contrast, for example, commentators attribute Apple's failure during the early 1990s to its refusal to give up old models of business success. Legacy models hindered Apple's development; the refusal to be captured by legacy models was a key to Microsoft's success.

40. In an environment where a company has power over the competitive environment itself, however, the rational incentives of a business may be different. If the business, for example, has control over the architecture of that competitive environment, then it will often have an incentive to design that architecture to better enable its legacy business models. As Charles R. Morris and Charles H. Ferguson describe it,

> Companies that control proprietary architectural standards have an advantage over other vendors. Since they control the architecture, they are usually better positioned to develop products that maximize its capabilities; by modifying the architecture, they can discipline competing product vendors. *In an open-systems era, the most consistently successful information technology companies will be the ones who manage to establish a proprietary architectural standard over a substantial competitive space and defend it against the assaults of both clones and rival architectural sponsors.*[8]

A company in this position can and will resist change, in order to keep doing what it knows best.

41. This was the problem with the telephone company prior to its break up by the government. The telephone monopoly enjoyed substantial returns from its existing network architecture. The fear of regulators was that these returns would make it unwilling to experiment with other architectures that might better serve communication needs. As we have said, there is at least some evidence that AT&T in fact resisted the emergence of the Internet because it feared its effect on AT&T's own business model. Certainly it resisted the development and interconnection of other technologies to its telephone network. The regulators who pushed to free the telecommunication network believed that the market would choose differently from how AT&T, as controller of the network, would choose.

42. Time has proven these regulators correct. Once freed from the strategic control of an entity that had a particular business plan to protect, the communications network has evolved dramatically. The competitive process was enabled by making the network neutral about its uses, and by giving competitors access to the network so that they could compete about its best use. The same wires that AT&T used to send analog voice only are now being used to deliver stock

---

[8] Charles R. Morris and Charles H. Ferguson, *How Architecture Wins Technology Wars*, **Harvard Business Review** 86, 88 (March April 1993) (emphasis added).

quotes, music, fantasy games, reference information — in short, the whole content of the Internet.

43. The lesson from this explosion of innovation is critically important. An architecture that maximizes the opportunity for innovation maximizes innovation. An architecture that creates powerful strategic actors with control over the network and what can connect to it threatens innovation. No doubt these strategic actors *might* behave in a pro-competitive manner. There is no guarantee that they will interfere to stifle innovation. But without competition or regulation to restrict them, we should not assume that they *will* somehow decide to act in the public interest.

III. The Proposed Merger of AT&T / MediaOne

A. The Threat Posed by Bundling

44. As we stated at the start, we do not question the merits of a merger between AT&T and MediaOne in principle. AT&T's argument that such a merger will enable much greater competition in local telephony may prove persuasive; the efficiencies of a merger for the supply of broadband access may prove persuasive as well. Our sole concern is the architecture that AT&T and MediaOne propose for broadband access. As they have described in their papers, cable broadband will prevent users from selecting an Internet Service Provider ("ISP") of their choice. Instead, access will be technologically bundled with ISP service. The network will thus discriminate in the choice of services that it allows. This kind of discrimination may have profound consequences for the competitive future of the net.

45. To see the problem with this architecture, we must first understand the importance of an ISP. ISPs serve a number of functions in the existing narrowband residential market. Some ISPs focus primarily on access to the Internet. Customers, through their local telephone

exchange, connect to the ISP; the ISP serves Internet access at speeds limited only by the local telephone exchange. Some ISPs supplement this access with promises of user support — both the support to assure the Internet connection is maintained, and in some cases, support with the use of certain Internet applications. Some ISPs further supplement this access with server capabilities — giving users the ability to build web pages on the ISPs servers, or support more expansive Internet activities. And finally some ISPs provide, or bundle, content with access. The most famous of these is America Online, but other ISP/content providers have included CompuServe and Prodigy.

46. This existing narrowband residential market is extraordinary competitive. Customers have a wide range of needs that customers have in this market; the market responds to this range of needs with different packages of services. Nationwide there are some 6,000 ISPs. In any particular geographic region, there can be hundreds that compete to provide service.

47. The functions performed by ISPs, however, are not fixed. They have no inherent "nature." Hence as bandwidth changes from narrow to broadband, we should expect the range of services offered by ISPs to change. As throughput becomes more critical in video services, for example, we could imagine ISPs competing based on the caching services they would offer. Or as the character of the content available increased, we might imagine some ISPs catering to certain content (video content) while others specialized elsewhere (new users).

48. The functions of ISPs, then, must not be conceived of too narrowly. Their importance, for example, has little to do with hosting "home pages" on the World Wide Web, or the portal sites they might now provide. Their importance is in the range of services they might bundle and offer competitively — from content (including video and audio services) to help